IMAGE DEBLURRING WITH
A SYSTOLIC ARRAY PROCESSOR

FIELD OF THE INVENTION

5

The present invention relates in general to image deblurring and, more particularly, to a method and device which includes an iterative update of the deblurred image estimate,

BACKGROUND OF THE INVENTION

Image distortions including noise contamination and blurring are encountered in many imaging applications. The earliest and most advanced applications are associated with astronomical imaging, space earth observations, airborne imaging, and forensic imaging. Deblurring is also a necessary part if the computed tomography imaging algorithms. The blurring might be caused by the imaging system being off-focus, or by the atmospheric distortions. In computer tomography the 3-D image restored through an inverse Radon transform is blurred and need to undergo deblurring. In recent years, the deblurring applications have proliferated into customer imaging devices as well, were they can be used for enhancement of static images.

The existing deblurring algorithms and methods require significant computing power and deblurring of high resolution image would take some time even for modern powerful computers. While this processing time delay is often acceptable for deblurring static images, such as astronomical images, it makes it difficult or impossible to enhance streaming images —such as digital video— in real-time, at the frame update rate.

A blurred image is usually considered as a transformation of an ideal image that needs to be restored. This transformation could be described as adding a noise and convolving the ideal image with a Point Spread Function (PSF) of the blur. The PSF describes a spatial intensity pattern that would be observed in the blurred image for a point source in

the ideal image. Herein, it is assumed that the PSF is known and spatially invariant ñ the same across the image.

There are two main classes of approaches to computing the deblurred image given the blur PSF. First class are the frequency domain approaches such as Wiener filter or regularized inversion in frequency domain. These approaches are described in more detail below and entail computation of a Fourier (Cosine) transform of the deblurred image. Applying the inverse Fourier transform then yields the sought deblurred image. The second class includes iterative updates approaches. These include Lucy-Richardson and other updates that iteratively optimize loss indices based on the image statistics. The commonly used iterative update methods are usually localized, in the sense that information from neighborhood pixels only is used to update the deblurred image estimate for a current pixel. The localization enables a parallelized systolic array implementation of the algorithm.

As mentioned above these existing methods work well for static images but are not well suited for real-time deblurring of streaming video. None of them is capable of addressing the entire list and finding a reasonable tradeoff between optimal restoration quality, noise insensitivity, localization, and computational performance.

Some prior art proposals for work with blurred images were developed in the early years of image processing and are very economical in terms of the required computing power and memory access. All of the three methods mentioned below include localized updates. An update for each pixel is based on the data for the pixels in

the immediate neighborhood. Such updates can be efficiently implemented through parallel processing using systolic arrays. Yet, as will be described below, the known iterative methods have a fundamental deficiency which is addressed by the proposed update.

5   The three iterative updates described below are most often suggested for use. The first iterative method is the Successive Approximation (Landweber) Method that corresponds to a steepest descent optimization of the quadratic error. It is also known as an iteration with reblurring and has the form:

10 $$u(n+1) = u(n) - \gamma H^*(H\, u(n)) - y_b$$

where n is the iteration number, γ is the scalar factor used to adjust the convergence speed and H^* is an adjunct operator for H. If H corresponds to a convolution kernel $h(-k,l)$ where $*$ denotes a complex conjugate.

15   A version of the update is the Van Clittered method which corresponds to the stochastic approximation or least-mean-square update minimizing the same quadratic performance index.

$$u(n+1) = u(n) - \gamma(H^* u(n)) - y_b$$

Both of these updates are linear. A nonlinear update that is 20 supposed to converge faster is the Lucy-Richardson update. While the first shown update above can be considered as a Maximum Likelihood optimization with Gaussian noise model, the Lucy-Richardson update is a Maximum Likelihood optimization with a Poisson noise model.

25

Assuming that H is normalized to unity, the Lucy-Richardson update can be compactly written in the form

$$u(n+1) = \left[\frac{y_b}{H * u(n)} * H' \right]$$

where the division or fraction inside the big brackets and the multiplication outside of the brackets are pixel-wise operations. The convolutions denoted by * are computed in the usual way.

These updates have two key features relevant to for this invention. First, computations in the above equations are localized to the extent that blur PSF operator H is localized. Also, the updates suffer from common problems including slow convergence noise amplification and 'ringing' (producing edge artifacts and spurious "sources"). For practical implementation of these updates, stopping the update in time is of foremost importance. The quality of the recovered image is first improved in the update and then starts deteriorating. Stopping an update in time to achieve optimal quality of the recovered image requires supervision and is unacceptable for an embedded implementation in a systolic array processor.

The proposed invention addresses the iterative update convergence issue and, thus, enables an embedded systolic array implementation. The main reason for the eventual divergence of the updates above is that each of these updates converges to the inverse solution as its steady state. The inverse solution is unacceptable because of high frequency noise amplification. To achieve good quality of deblurring, there is a need for high-frequency regularization, an optimal trade off between idealized restoration accuracy and noise amplification.

deblurring, there is a need for high-frequency regularization, an optimal trade off between idealized restoration accuracy and noise amplification.

It would be of great advantage if an iterative update would provide high-frequency regularization.

5   Another advantage would be if an iterative update would provide an optimal trade off between idealized restoration accuracy and noise amplification.

Other advantages and features will appear hereinafter.

SUMMARY OF THE INVENTION

This invention proposes a method and device for image deblurring. The proposed method is an iterative update of the deblurred image estimate. The update has two principal terms: (i) feedback of blurred image prediction error using the deblurred image and (ii) feedback of the past deblurred image estimate added to the discrete integrator. The two feedback terms are based on two feedback operators, which are localized FIR convolution operators.

The update feedback operators can de designed with a given degree of localization to provide an optimal tradeoff solution taking into account multiple design objectives including update convergence speed, noise amplification, quality of image restoration, and robustness to operator implementation error. The solution to the operator design problem is computed by Linear Programming (LP) optimization. The update, using the local feedback operators is implemented by using a systolic array processor, where each process can be mass-produced and is capable of simple multiplication and addition operations, as well as communication with the neighbors. With the localized operators, a parallelized implementation of the update algorithms is possible that is completely scalable to very large image sizes. Since the feedback operators are localized, the data exchange necessary to perform the update computations is limited to a neighborhood of each array node processor and computations do not depend on the image size. The proposed update being implemented on an inexpensive systolic processor enables real-time deblurring of high-resolution video images.

BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the invention, reference is hereby made to the drawings, in which:

5   FIGURE 1 is a schematic illustration of a systolic array processor;

FIGURE 2 is a plan view of the processor of FIGURE 1, showing interconnecting logic blocks;

FIGURE 3 is an illustration of an original undistorted image;

FIGURE 4 is an illustration of a Gaussian blur operator applied to

10   the image of FIGURE 3;

FIGURE 5 is an illustration of the optical transfer function of the blur operator of FIGURE 4;

FIGURE 6 is a an illustration of a blurred image of FIGURE 3;

FIGURE 7 is a an illustration of a designed symmetric update

15   operator K;

FIGURE 8 is a an illustration of a designed symmetric smoothing operator S and

FIGURE 9 is a an illustration of the recovered or deblurred image.

In the figures, like reference characters designate identical or

20   corresponding components and units throughout the several views.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

In its simplest form, the present invention provides a method and device for deblurring an image having pixels. A blurred image is

5 downloaded into a systolic array processor having an array of processing logic blocks such that each pixel arrives in a respective processing logic block of one pixel or small groups of pixels. Data is sequentially exchanged between processing logic blocks by interconnecting each processing logic block with a predefined number of adjacent processing

10 logic blocks, followed by uploading the deblurred image. For example if a 1024 by 1024 pixel image is selected, the array may be 256 by 256, thus processing 4 by 4 groups of pixels. Other groups including single pixels, 2 by 2 groups, 3 by 3 groups, and so on may also be processed by the selected systolic array processor.

15 The processing logic blocks provide an iterative update of the blurred image implemented in the processing logic blocks by $u(n + 1) =$ $u(n) - K^*(H^* u(n) - y_b) - S^* u(n)$ where u is the ideal undistorted image, m and n are column and row indices of an image pixel element, $y_b (m,n)$ is the observed blurred image, * denotes a 2-D convolution, K is

20 a feedback update operator with a convolution kernel $K(m,n)$ and S is a smoothing operator with a convolution kernel $S(m,n)$.

It is important to distinguish between deblurring algorithm design, which is usually a result or subject of mathematical analysis and its hardware implementation that allows optimized computational

25 performance. In this invention, these two steps are brought closely

together such that the very computational structure of algorithms is such as to allow high performance implementation. At the same time, the mathematical analysis and design explicitly take this structure as one of the design constraints.

5   Referring to Fig. 1, data flow in the distributed update $u(n+1) = u(n) - K^*(H^*u(n) - y_b) - S^* u(n)$ shows how each pixel, the intermediate data in the update is stored in three planes: the blurred image y_b, the current deblurred image estimate u and the prediction error $H^*u - y_b$. The update uses data from a localized neighborhood of

10   each pixel in each of the three planes.

The update $u(n+1) = u(n) - K^*(H^*u(n) - y_b) - S^* u(n)$ is amenable to a systolic array processor implementation. Such processor illustrated in Figure 2 consists of an array of simple processing elements (logic blocks), one per image pixel. The processing logic blocks are

15   interconnected such that each can exchange data with its immediate neighbors. Each of the processing logic blocks also has some local data memory and is capable of simple arithmetic operations such as addition, subtraction, and multiplication. Using the interconnect logic, the blurred image can be downloaded into the processor array, each pixel

20   into a respective processing logic block. By sequentially exchanging data with the nearest neighbors, each logic block can accumulate data from any predefined number of neighbors within certain reach. Of course this would require multiple update cycles, a larger number of cycles for a larger reach.

By accumulating the data within the reach demanded by the localized spatial operators in the update above for each pixel can be implemented within the respective processor. This presumes preloading and storing the operators H, K and S in each of the array
5   processing logic blocks. For each pixel, the computations are extremely simple and require a number of additions, subtractions, and multiplications, proportional to the size (squared reach) of the FIR operators H, K and S. Apart from the data transfer associated with downloading the blurred image and uploading the deblurred image
10   estimate, the computational demand on each of the processing logic blocks does not at all depend on the image size. Of course, the overall amount of computations grows linearly with the image size. These computations are, however, performed in parallel by the processing logic blocks of the array. For large array size, an overall computational
15   power of the described specialized parallel computer implemented by the systolic array could be very substantial. The data transfer requirements grow very slowly, as a square root of the image pixel size. Thus, the proposed approach is fully scalable and can be implemented at a very high speed for high resolution video images. Even for inexpensive
20   systolic array processor, it is possible to do a few dozen update iterations within a video frame update time.

The update above can be considered implementing a spatial 2-D IIR filter. This 2-D IIR filter allows satisfying design requirements much better than a FIR filter with the same number of spatial tap delays. The
25   above described use of a systolic array processor to implement the 2-D

IIR filter can be considered as an extension of well-known practice of using shift registers and summer/adder logic in high-performance hardware implementations of standard time-domain digital filters. The additional complexity in this invention is caused by dealing with 2-D

5   signals (images) that are iterated in time.

The design of localized FIR operators K and S herein closely follows the approach of [8], where feedback design for a distributed control problem is considered. The update $u(n + 1) = u(n) - K^*(H^* u(n)$ $- y_b) - S^* u(n)$ is a recursive filter estimating the 2-D input signal u

10   from the noisy blurring system output data. It is well known that estimation problems in linear systems theory are dual to control problems and the same mathematical tools can be applied to both.

The FIR feedback operators K and S in $u(n + 1) = u(n) - K^*(H$ $^* u(n) - y_b) - S^* u(n)$ are assumed to have the same symmetry

15   properties as the blur operator H. In the basic case of central (2-fold) symmetry, these operators can be presented in the form

$$K(\lambda_1 \lambda_2) = k_o + \sum K_{mn}(\lambda^{m}_1\lambda^{n}_2 + \lambda^{-m}_1\lambda^{-n}_2),$$

$$S(\lambda_1 \lambda_2) = s_o + \sum S_{mn}(\lambda^{m}_1 \lambda^{n}_1 + \lambda^{-m}_1 \lambda^{-n}_1),$$

where k_o, K_{mn}, s_o, s_{mn} are the design parameters and the summation in

20   these equations is performed over all indices covering a given vicinity of the zero spatial shift. A rectangular support set of the FIR operators might consists of all indices such that $0 < m < N$, $|n| < N$ or $m = 0$, $0 < m < N$.

25

Respective relationships can be written down in the case of a 4-fold or 8-fold symmetry. As an example, in case of a 8-fold symmetry the operator K has the form

$$K(\lambda_1\ \lambda_2) = K_o + \sum K_{mn}(\ \lambda^m_1\lambda^n_2 + \lambda^{-m}_1 + \lambda^{-n}_2 + \lambda^{-m}_1\lambda^n_2 + \lambda^m_1 + \lambda^{-n}_2 + \lambda^n_1\lambda^m_2 + \lambda^{-n}_1$$

5
$$+\lambda^{-m}_2 + \lambda^{-n}_1\lambda^m_2 + \lambda v_1\lambda^{-m}_2),$$

Collecting all the independent weights of the two FIR operators yields the design parameter vector

$$x = [K_o\ ...\ K_o\ ...\ K_{mn}\ ...\ S_o\ ...\ S_{mn}\ ...]$$

This vector should be chosen such that the system transfer

10   functions satisfy requirements the above equations.

By substituting $\lambda_1 = e^{jv_1}$, $\lambda_2 = e^{jv_2}$ into the above equations, one can notice that all three optical transfer functions $K^\sim(v_1\ v_s)$ $S^\sim(v_1\ v_s)$ and $H^\sim(v_1\ v_s)$ are real. The functions $K^\sim(v_1\ v_s)$ $S^\sim(v_1\ v_s)$ and $K^\sim(v_1\ v_s)H^\sim(v_1\ v_s)$ are also linear in the design parameter vector. Another key fact is

15   that the denominator is always real positive. Being affine in x this denominator can be presented in the form

$$S^\sim(v_1\ v_s) + K^\sim(v_1\ v_s)H^\sim(v_1\ v_s) = c^T(v_1\ v_s)x$$

$$c^T(v_1\ v_s)x\ 0,\quad v_1\ v_s)\in\Re$$

The last inequality follows from the requirement which can be

20   presented in the form

$$1 - r_0 \le c^T(v_1\ v_s)x \le 1 + r_0$$

The design constraints and other design specifications that one might wish to consider for this problem can be presented in the form

25
$$\left|\ \frac{c^T(v_1\ v_s)x + b(v_1\ v_s)x}{c^T(v_1\ v_s)x}\ \right|$$

By using this denominator positively, this can be written as

$$-c^T(v_1,\ v_s)x \le d^T(v_1,\ v_s)x + b(v_1,\ v_s) \le c^T(v_1,\ v_s)x$$

For controller design, these inequalities can be complemented by a requirement that the integrator leakage operator S is possibly small.

5   Small operator S improves the image recovery performance. The latter requirement can be presented in the form

$$|R_x| \le x_0, \quad x_0 \to \min.$$

where the matrix R selects the last half of the components of x, ones that contain the weights of the smoothing operator S. The absolute value

10 above is component-wise and x_0 is a scalar. The linear inequalities of the form of the denominator following from the design constraints as well as can be solved by gridding the spatial frequencies { x_1, x_2}. By adding the requirement and augmenting the design vector x with the parameter x_0 we have a Linear Programming (LP) problem. The LP problem even of a

15 very large size (for a dense grid in v_1, v_2) can be solved efficiently and reliably with help of modern interior point solvers. In the example below a LINPROG solver from Matlab Optimization Toolbox was used.

In the event that the image that is to be deblurred is a color image, the update equation is slightly different. Specifically, for a color image,

20 $y_b = y_d(j,k,c)$ where c denotes the color. For instance, the color mmay be 'r', 'g' or 'b' and $y_b = y_d(j,k,c)$ is the intensity of the color c at the pixel with the coordinates (j,k). For an hyperspectral image, there might be many more thhan three 'colors' ir waavelengths in the image. For a color image, the update has the same form as above for each color

25 and can be written in the form such that the iterative update is

implemented in the processing logic blocks by $u(n+1; c) = u(n, c) - K$ $*(H^* u(n, c) - y_b(c)) - S^* u(n, c)$ where $y_b(c) = y_d(j, k, c)$ is the 2-D array of color c intensities for the blurred image encompassing all pixels (j, k) in the image and $u(n, c) = u(j, k, n, c)$ is the 2-D array of color c intensities for the restored image estimates at iteration number n. The image $u(n, c)$ is an array that includes all pixels (j, k).

EXAMPLE AND ILLUSTRATION

The image in Figure 3 was distorted with a localized Gaussian blur operator H with a PSF illustrated in Figure 4. This PSF is a FIR operator with maximal $N=3$ spatial offset taps on each side for each of the two spatial dimensions. The optical transfer function $h(v_1, v_s)$ of the blur in Figure 3 is displayed in Figure 5. The original image of Figure 3 was blurred with the operator H in Figure 4 and a random noise e with a maximum magnitude 0.02 was added. The blurred and noisy image y_b is shown in Figure 6.

The feedback operators S and K have been designed as described in the previous section. The design assumed FIR operators with $N=3$ maximal spatial offset and 8-fold symmetry (the operator H has 8-fold symmetry). The following parameters have been used in the design:

- the convergence rate was chosen as $r_0 = 0.08$
- the recovery error bound was chosen as $u_0 = 0.3$
- the image noise bound was chosen as $d_0 = 0.08$

- the in-band domain B was chosen as a set of spatial frequencies such that $H^{-}(v_1\ v_s) \leq h_0$, where $h_0 = 0.185$

The design resulted in the feedback operators K and S illustrated in Figures 7 and 8 respectively. Because of the 8-fold symmetry, only a
5 first quadrant (nonnegative offsets) is shown.

The recovered image was obtained by applying 30 steps of the update with designed operators K and S and the known blur operator H to the image y_b shown in Figure 6. The recovered image is displayed in Figure 9. As one can see the recovery quality is quite good. The original
10 and blurred image were taken from Matlab Image Processing Toolbox demo illustrating the use of a Lucy-Richardson algorithm. The image in Figure 9 has somewhat better recovery quality than one obtained by the Lucy-Richardson algorithm in the Matlab demo. The fundamental difference however is that run further unchecked, the update continues
15 (slightly) improving the recovery quality, while continuing the Lucy-Richardson update leads to the restored image deterioration because of the high-frequency noise increase and ringing that become increasingly noticeable.

This invention makes possible fast real-time deblurring with help
20 of inexpensive hardware. This enables a host of applications with customer-grade imaging devices such as: scanners, photocopiers, cameras, and various streaming video devices. As one example, security video or other remote camera video might be slightly off focus with no possibility of focusing the remote camera. This video could be enhanced
25 by deblurring.

While particular embodiments of the present invention have been illustrated and described, they are merely exemplary and a person skilled in the art may make variations and modifications to the embodiments described herein without departing from the spirit and scope of the present invention. All such equivalent variations and modifications are intended to be included within the scope of this invention, and it is not intended to limit the invention, except as defined by the following claims.